Real Users – Real Requirements

Jennifer M. Schopf
UK National eScience Centre,
Argonne National Laboratory

Steven J. Newhouse
Open Middleware Infrastructure Institute

Abstract

Pervasive Grid adoption is predicated on the availability of widely deployed usable software and a user community willing to use it. Currently, widespread adoption of Grids, even within technically sophisticated communities, is extremely limited. Determining and eliminating these barriers to adoption are essential in order for Grids to becoming widely adopted. Through a series of face-to-face interviews conducted during summer 2004, we have identified issues relating to job submission, file transfer, usability, and systems management that must be resolved in order to improve the usability of grid infrastructures. The background to these issues and some possible solutions are described in this paper.

1 Background and Motivation

During July and August 2004 we visited various applied science and middleware groups in the U.K. in order to gather basic information on the services and functionality these projects are using. Our motivation was to help guide the development of future activities and priorities within the U.K.'s Open Middleware Infrastructure Institute [OMII] and the Globus Alliance [Globus] and to inform the wider Grid community of the status of some current services. To understand their issues in more detail, we held meetings with application developers with some Grid (generally Globus Toolkit 2 or Globus Toolkit 3) or Web services experience, those with software that might be of broader use or interest, and those who have expressed dissatisfaction with current tools.

The twenty-five groups, listed below, included representative applications from biology, chemistry, physics, climatology, and other scientific fields, as well as a smaller set of basic tool builders. In addition, informal discussions took place at several workshops that allowed us to obtain a broader perspective as to the important issues in the community.

We asked what functionality the groups had tried in the past, what their applications needed today from the current Grid infrastructures, and what functionality the group was considering for near-future plans. Most meetings ended by our inquiring what functionality the group thought was most important and still lacking in today's tools or services.

2 Topics of Concern

Through our initial meetings we identified a set of open areas of concern that were repeated by many groups.

Almost every group we spoke with was using Grid environments to support applied science activities. The functionality the group members wanted was for their day-to-day work, not farther-out speculative needs. These primarily fell into two categories: job submission and tracking, and file transfers. A few projects were using tool "add-ons," such as visualization tools, data format translators, or policy management tools, but these were always strongly tied to the project domain and narrowly scoped. When we asked about other possible functionality or services that could be used, we were told these were not on the six-month horizon most groups were currently interested in, as detailed below.

2.1 Job Submission

Job submission for most projects included a simple, dependable, "run my application" interface that was in the users "comfort zone" and behaved as expected. Many users had adapted a standard tool, such as the Globus Toolkit job submission [GRAM] or Condor [Condor], for project-specific use. Job submission was generally being done on well-known resources or services.

Having conquered the initial challenge of job submission using Grid tools, users are now concerned with understanding where a job is in its lifetime, where it is failing, why, and what to do next. There was an expressed concern with the lack of tools to understand failures or performance faults in general.

With respect to resource discovery, most projects were using a small set of well-known resources or services, so this simply was not an issue. However, when a choice of resources was available, their status was found through manually checking as opposed to automated tooling.

2.2 File Transfer

Most users were transferring files by using Grid tools such as GridFTP and were happy with the service level they experienced. Some groups needed reliable file transfers, either because they had many small files to transfer and it was easy for one in a thousand to have problems and be left behind, or because they had such large files that the file transfer time was greater than the mean time to failure for some system component, often an unreliable network connection.

A few groups we spoke with were beginning to examine higher-level file transfer services, such as provenance services, access to databases, or replication, but these groups were still primarily prototyping these efforts. This may be in part because of the groups we interviewed – very pragmatic, production applications, as opposed to those in the development phase.

The tool most commonly requested with respect to file transfers was one to help diagnose the problems, including that of slow performance, seen on systems when performing large file transfers. Invaluable would be a tool to help users understand where a problem is being caused so they can better understand who to contact.

2.3 Usability Issues

All of the groups we spoke with stated that if tools weren't within a user's comfort zone, they simply were not considered for use.

Most users want a layer between them and the tool in order to bring the functionality into their own comfort zone. These wrappers do not add functionality per se but do significantly increase the usability and usefulness of a service.

One suggestion by several groups was to have small tools each of which solved a single function that could then be composed together. These chains of services would allow a number of different use cases based on the same set of basic tools. However, there was a tension between the need for these to meet end-to-end application specific needs and yet have them be generic enough to be used between projects.

As mentioned above, the lack of user-oriented diagnostic tools was a significant problem. Most diagnostic tools solve problems other than those seen at the user-level. Tools that look like normal user applications and can help an average user diagnose failures are a strong current need.

Another continuing concern was the need for training, especially for security concerns. Security is seen as extremely challenging; and system administrators, developers, and users all want more information about common practices and current approaches.

2.4 System Administration Concerns

The primary system administration concern that we encountered was that of dependable builds. Software that didn't build dependably was seen as simply not under consideration for further use – no matter what the expected functionality of the tool. Too often it was found that Grid software was built nondeterministically or did not have installation verification tools available.

Verification in general was seen as a significant problem. With the overall time to failure for Grid components decreasing as their number increases, there is a strong need for better verification and instability analysis to discover and resolve problems before a user happens upon them.

3 Conclusion

Over the course of several weeks in July and August 2004 we spoke with 25 UK eScience project groups about their use of Grid functionality and services. What resulted is a picture of current application and user needs of these services, and some suggestions for ways to move forward. This data is now influencing the directions of both the Globus Alliance and the OMII.

The strongest result that came from these discussions was the simple need for on-going conversations between tool developers and users. Grid tool developers must continue to talk and interact with application scientists; without such interaction, the tools are for nothing.

Both organizations feel that they have benefited from this activity and that it has helped in understanding the expectations and requirements of the user community.

Acknowledgments

This work has been supported by the UK e-Science Core programme (through the Open Middleware Infrastructure Institute) and by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-ENG-38. We thank all the people we visited for their hospitality and the uninhibited technical discussions.

References:

[Condor] "The Condor Project," http://www.cs.wisc.edu/condor/

[Globus] "The Globus Alliance," http://www.globus.org

[GRAM] "Grid Resource Allocation and Management (GRAM)," http://www.globus.org/gram

[OMII] "Open Middleware Infrastructure Institute," http://www.omii.ac.uk

Projects

The projects we interacted with as part of this work included the following:

M. Baker, Portsmouth, OGSA Testbed
http://dsg.port.ac.uk/projects/ogsa-testbed

R. Baldock, MRC Human Genetics Unit and NeSC, Edinburgh Mouse Atlas Project
http://genex.hgu.mrc.ac.uk/

R.Baxter, EPCC, eDIKT
http://www.edikt.org/

D. Chadwick, Salford, PERMIS
http://www.permis.org/

N. Chue Hong, EPCC, OGSA-DAI
http://www.ogsadai.org.uk/

D. Colling, IC, GridPP2
http://www.gridpp.ac.uk/

T. Cooper-Chadwick, Southampton, gYacht/gShip
http://www.soton.ac.uk/~gyacht/

S. Cox, Southampton, GeoDise
http://www.geodise.org/

M. Daw, Manchester, Access Grid & MUST
http://www.agsc.ja.net
http://www.sve.man.ac.uk/Research/AtoZ/MUST/

W. Emmerich, UCL, eMinerals and OGSI Testbed
http://eminerals.org/

M. Ghanen, Imperial, DiscoveryNet
http://www.discovery-on-the.net/

M. Giles, Oxford, gViz
http://www.visualization.leeds.ac.uk/gViz/

C. Goble & N. Sharmen, Manchester, myGrid and Integrative Biology Project
http://www.mygrid.org.uk/
http://www.integrativebiology.ox.ac.uk/

S. Lloyd, Oxford, eDiamond
http://www.ediamond.ox.ac.uk/

J. MacLaren and J. Brooke, Manchester, Brokering activities at Manchester Computing
http://uombroker.sourceforge.net/docs/server/overview-summary.html

A. Martin, Oxford, ClimatePrediction.NET
http://climateprediction.net/

M. McKeown, Manchester, OGSI:Lite and WSRF:Lite
http://www.sve.man.ac.uk/Research/AtoZ/ILCT

Andy McNab, Manchester, GridPP2
http://www.gridpp.ac.uk/

S. Pickles, Manchester, TeraGyroid and GRENADE
http://www.realitygrid.org/TeraGyroid.html
http://mrccs.man.ac.uk/research/grenade

A. Porter, Manchester, RealityGrid, and M. Rider, Manchester, eViz
http://www.eviz.org/

A. Rector, Manchester, CLEF
http://www.clinical-escience.org/

R. Sinnott, Glasgow, BRIDGES
http://www.brc.dcs.gla.ac.uk/projects/bridges/public/people.htm

T. Sloan, EPCC, INWA
http://www.epcc.ed.ac.uk/inwa/

L. Smith, EPCC, QCDGrid
http://www.epcc.ed.ac.uk/computing/research_activities/grid/qcdgrid/

L. Yang, B. Yang, NeSC, AI Workflow
http://dream.dai.ed.ac.uk/e-Science/